Reducing energy consumption and carbon footprint of data centers through heterogeneous leverage management

Eco-ICT 2024

Vladimir Ostapenco [1]
Laurent Lefevre [1]
Anne-Cécile Orgerie [2]
Benjamin Fichel [3]

[1] Univ Lyon, EnsL, UCBL, CNRS, Inria, LIP
[2] Univ Rennes, Inria, CNRS, IRISA
[3] OVHcloud

Context

- Significant increase in the use of information and communication technologies (ICT)
- GHG emissions related to ICT were estimated to be between 2.1% and 3.9% of global GHG emissions in 2020 [1]
  - Risk of almost doubling and reaching more than 7% by 2025 [2]
- Data centers are one of the essential components of ICT
  - Impacts are dominated by the usage phase, containing mostly electricity consumption [3]
  - Estimated to consume approximately 1% of global electricity [4]

[1] Freitag et al. (2020)
[2] The Shift Project (2021)
[3] Malmodin et al. (2024)
[4] Masanet et al. (2020)

Motivation

- ***Numerous approaches*** (energy/power leverages) ***have been proposed*** to reduce the energy usage of data centers [1 2 3]
- ***Multi-criteria impacts must be considered***
  - Covering GHG emissions, pollution, and impact on air, water, and biodiversity
- We define ***heterogenous leverage*** as a leverage that can ***cover both power/energy consumption and carbon footprint*** of data centers
- Heterogeneous leverages
  - ***Technological*** (e.g. powering on/off or slowing down resources (DVFS, RAPL))
  - ***Logistical*** (e.g installation and decommissioning of machines)
- ***Challenge*:** Manage these heterogeneous leverages in a large-scale environment
  - To improve the overall energy and carbon footprint of the entire datacenter infrastructure

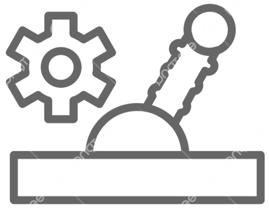

[1] Raïs (2018)
[2] Ayanoglu (2019)
[3] Orgerie et al (2014)

Heterogeneous leverage modeling

Types of leverages

Atomic action-based leverage

- Performs a directly applicable action on a single data center component

Examples :

- Power limiting of a compute node based on Intel RAPL technology
- Compute node deployment

Intelligent action-based leverage

- Performs a set of combined actions on multiple data center components
- Creates atomic action-based leverages on one or more datacenter components

Examples :

- Power capping of an entire data center
- Compute node end-of-life management

Evaluation of leverages for efficient leverage management

- Each leverage has
  - Various **impacts** on data center components (e.g. power draw, performance)
  - **Costs of applying** in terms of energy, time and even carbon footprint
- Assessment of costs and impacts using all possible combinations of leverage states is required
  - ***Results in a table of leverages***

Table of leverages of DVFS leverage

States (Frequency)	Impacts (Power)	Impacts (Performance)	Costs (Time)
1.2 Ghz	17 W	0.4	1ms
1.4 Ghz	18 W	0.6	1ms
1.8 Ghz	22 W	0.7	1ms
2.0 Ghz	24 W	0.9	1ms
2.6 Ghz	26 W	1	1ms

Table of leverages of compute node deployment leverage

States	Impacts (Rack space units)	Costs (Time)	Costs (Carbon footprint)
Initiated	0	0 days	0 kgCO_2eq
In-Progress	0	0 days	0 kgCO_2eq
Finished	1	24 days	1306 kgCO_2eq

Dell PowerEdge R640

Applying leverages to data center components: Environmental Gantt chart

- Timeline of the complete life cycle of each data center component
- Contains leverages that have been applied to each data center component since the beginning of its life

Simplified example of environmental Gantt chart

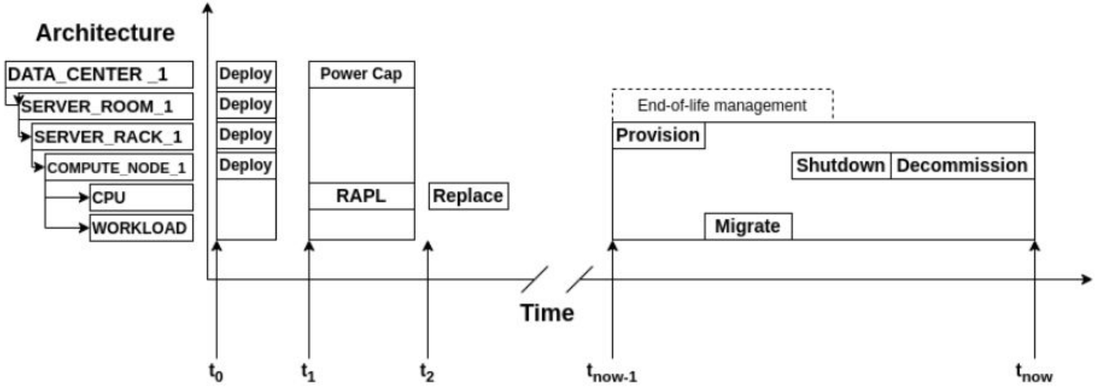

Usage scenarios

- Leverage positioning for
  - Instant infrastructure management
  - Planning of actions in the future
- Metrics tracking
  - Energy consumption
  - Carbon emissions
- Replay of the past
- Dashboard
- Life Cycle Assessments
  - Follow the use and evolution of data center components

Management of leverages: Leverage management framework

- Management of leverages and their combinations
  - Analyzes infrastructure metrics and leverages impacts and costs
  - Chooses a set of leverages to apply
  - Places leverages on the environmental Gantt chart
- In order to
  - *Meet constraints (e.g.* power budget, cooling capacity)
  - *Reduce impacts* (e.g. energy consumption, carbon footprint)
  - *React to external events* (e.g. unfavorable weather conditons)
  - *Respond to a request of a cloud provider operator*
- Leverage execution monitoring
  - Detect and respond to problems when applying leverages

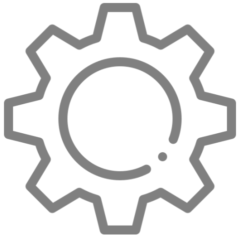

Validation: Temporary power capping at data center scale

Modeled, evaluated and implemented

- Atomic action-based leverages
  - Intel RAPL power limiting leverage
  - Shutdown leverage
- Intelligent action-based power capping leverage

Scenario operation

- Management framework places the intelligent action-based leverage with a power reduction objective on the Gantt Chart
- Intelligent action-based leverage finds a placement of atomic action-based leverages and places them on the Gantt Chart
- Leverage management framework verifies if power constraint is respected

Simulation

- Small data center with 5 compute nodes
- Power reduction objectives ranging from 10% to 90%

Power capping scenario execution environmental Gantt chart

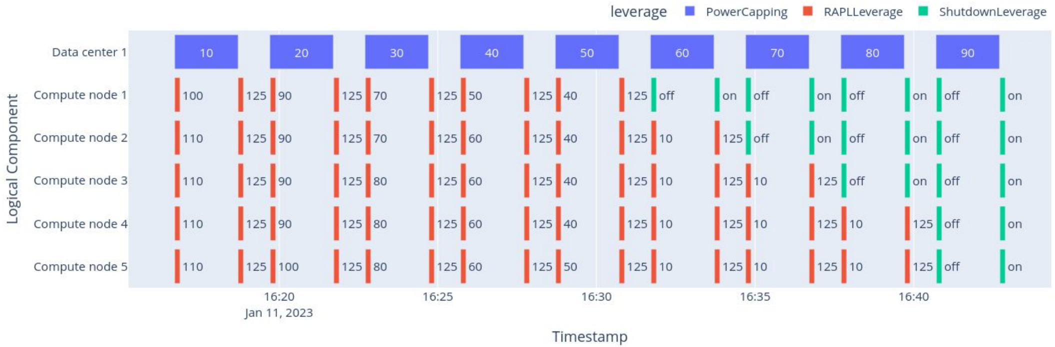

- Instantaneous data center consumption can be significantly reduced by only slowing down compute nodes without canceling client workloads

Validation: Minimizing the global carbon footprint of a data center

- Modeled and simulated the ***standard infrastructure management*** process used by a cloud provider
  - Increasing workload management leverages
  - Compute node decommissioning leverages
- Created, modeled and simulated a ***minimization scenario management*** process
  - Workload re-balance leverages
  - Idle compute nodes management leverages
- For both scenarios
  - ***Placed the leverages*** on the environmental Gantt Chart
  - ***Calculated the carbon footprint*** of the infrastructure over a 10-year period

Evolution of carbon footprint over a 10-year period showing a significant reduction in the carbon footprint for the minimization scenario

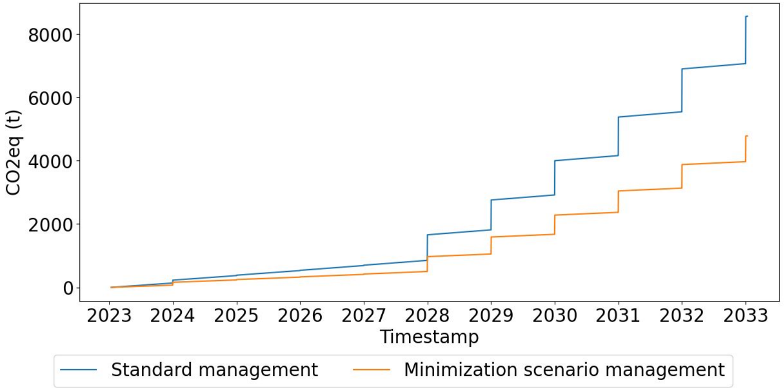

Our approach allows

- Impacts accounting and estimation
- Impacts forecasts
- Further optimization of the infrastructure management

Conclusion

- Data centers must enter in a **new era of reduced environmental impacts**
- Providers face **challenges in management** of technological and logistical capabilities (leverages)
- Proposed approach based on
    - Heterogeneous leverage modeling and evaluation
    - Environmental Gantt Chart
    - Leverage management framework
- Opens the door to some large-scale facilities to **integrate, plan, map, monitor and orchestrate environmental leverages**

For more details

- Ostapenco, Vladimir, Laurent Lefèvre, Anne-Cécile Orgerie, and Benjamin Fichel. **"Modeling, evaluating, and orchestrating heterogeneous environmental leverages for large-scale data center management."** *The International Journal of High Performance Computing Applications* 37 (May 2023)
- https://hal.science/hal-04047008